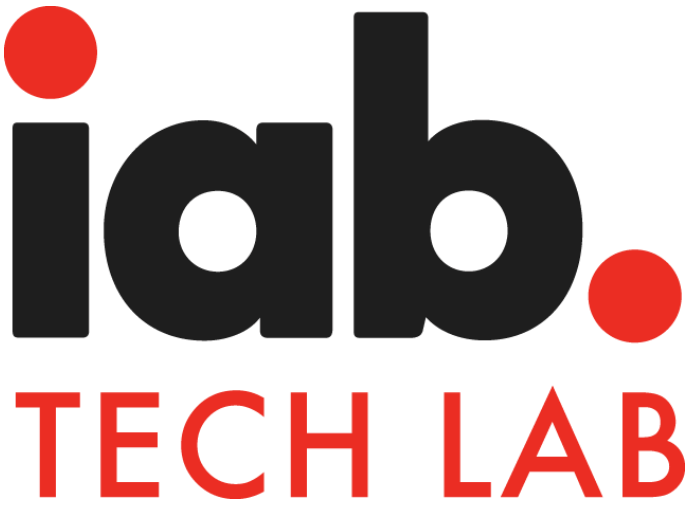

AI in Advertising Primer

Please email support@iabtechlab.com for questions, public comments and feedback. This document is available online here.

About this document

Artificial Intelligence is rapidly evolving the way that we think about everything we do. In the technology heavy field of digital advertising, we're no stranger to machine learning algorithms, they've been with us for over a decade. The rapid evolution of generative AI brings opportunities and challenges for us as technologists in digital advertising.

While this primer won't seek to answer every question conclusively. We seek to lay a foundation of terminology and outline some consistent use cases in the digital advertising industry.

This document is intended to be an informational guide to understanding Artificial Intelligence and functions to help organizations evaluate different AI use cases against their own brand and business requirements.

This document is developed by a subcommittee of the IAB Tech Lab Board of Directors. It seeks to evaluate the landscape and foster discussion on any necessary standards and specifications in the field.

Note: *The use of words or phrases 'Privacy", "Private", "Security", "Control", "Processing", "Personal Data", "PII" in this document is generic and does not refer to definitions in any specific regulation e.g. GDPR or CCPA.*
Throughout the document the word or phrases "ID", "user ID", "Consumer ID", are used interchangeably referring to a unique identifier associated with a user of a service.

License

Significant Contributors

This document was created by a IAB Tech Lab Board of Directors subcommittee, with significant contribution from, Michael Palmer, GroupM, Trey Griffin, Raptive, Aaron Brown, Madhive, David Caragliano, Google, Rebecca Gleason, Meta, Jordan Rogers-Smith, Meta

IAB Tech Lab Lead

Jill Wittkopp, VP Product

About IAB Tech Lab

The IAB Technology Laboratory is a nonprofit research and development consortium charged with producing and helping companies implement global industry technical standards and solutions. The goal of the Tech Lab is to reduce friction associated with the digital advertising and marketing supply chain while contributing to the safe growth of an industry.

The IAB Tech Lab spearheads the development of technical standards, creates and maintains a code library to assist in rapid, cost-effective implementation of IAB standards, and establishes a test platform for companies to evaluate the compatibility of their technology solutions with IAB standards, which for 18 years have been the foundation for interoperability and profitable growth in the digital advertising supply chain. Further details about the IAB Technology Lab can be found at https://iabtechlab.com.

Disclaimer

THE STANDARDS, THE SPECIFICATIONS, THE MEASUREMENT GUIDELINES, AND ANY OTHER MATERIALS OR SERVICES PROVIDED TO OR USED BY YOU HEREUNDER (THE "PRODUCTS AND SERVICES") ARE PROVIDED "AS IS" AND "AS AVAILABLE," AND IAB TECHNOLOGY LABORATORY, INC. ("TECH LAB") MAKES NO WARRANTY WITH RESPECT TO THE SAME AND HEREBY DISCLAIMS ANY AND ALL EXPRESS, IMPLIED, OR STATUTORY WARRANTIES, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AVAILABILITY, ERROR-FREE OR UNINTERRUPTED OPERATION, AND ANY WARRANTIES ARISING FROM A COURSE OF DEALING, COURSE OF PERFORMANCE, OR USAGE OF TRADE. TO THE EXTENT THAT TECH LAB MAY NOT AS A MATTER OF APPLICABLE LAW DISCLAIM ANY IMPLIED WARRANTY, THE SCOPE AND DURATION OF SUCH WARRANTY WILL BE THE MINIMUM PERMITTED UNDER SUCH LAW. THE PRODUCTS AND SERVICES DO NOT CONSTITUTE BUSINESS OR LEGAL ADVICE. TECH LAB DOES NOT WARRANT THAT THE PRODUCTS AND SERVICES PROVIDED TO OR USED BY YOU HEREUNDER SHALL CAUSE YOU AND/OR YOUR PRODUCTS OR SERVICES TO BE IN COMPLIANCE WITH ANY APPLICABLE LAWS, REGULATIONS, OR SELF-REGULATORY FRAMEWORKS, AND YOU ARE SOLELY RESPONSIBLE FOR COMPLIANCE WITH THE SAME, INCLUDING, BUT NOT LIMITED TO, DATA PROTECTION LAWS, SUCH AS THE PERSONAL INFORMATION PROTECTION AND ELECTRONIC DOCUMENTS ACT (CANADA), THE DATA PROTECTION DIRECTIVE (EU), THE E-PRIVACY DIRECTIVE (EU), THE GENERAL DATA PROTECTION REGULATION (EU), AND THE E-PRIVACY REGULATION (EU) AS AND WHEN THEY BECOME EFFECTIVE.

Table of Contents

Glossary

Artificial Intelligence (AI)	A simulation of human intelligence processes by machines, especially computer systems. Examples include decision making, visual perception, speech recognition and language translation.
Content Provenance	Provenance refers to the basic, trustworthy facts about the origins of a piece of digital content (image, video, audio recording, document). It may include information such as who created it and how, when, and where it was created or edited. The content author always has control over whether provenance data is included as well as what data is included. Included information can be removed in later edits. Provenance also allows for anonymous content.
Context Window	The maximum number of tokens (words or pieces of words) that a model can use at one time when generating media or making predictions.
Copyright	In the context of LLMs, copyright pertains to the legal rights associated with the content generated by these models. LLMs raise questions about who owns the rights to the output produced: the creators of the model, the users who prompt the responses, or the sources of the training data. This is a complex and evolving area of law, as LLMs have the potential to generate content that is derivative of or similar to copyrighted material included in their training datasets.
Deep Learning	A subset of machine learning that uses artificial neural networks with multiple layers to learn from data.

Diffusion Model	A diffusion model is a deep neural network that holds latent variables capable of learning the structure of a given image by removing its blur (i.e., noise). Diffusion models power image generators like Midjourney, Dall-E 3, and StableDiffusion. This is also the underlying technology behind text-to-video models like Pika, Runway Gen-2, and StableVideo.
Fair Use	A legal doctrine that permits limited use of copyrighted material without requiring permission from the rights holders. This concept is particularly relevant to LLMs due to the way they are trained and how they generate outputs. There is an evolving legal debate as to whether LLMs, image diffusers, and other generative AI systems qualify for Fair Use protection. Please note that Fair Use has a specific market connotation or jurisdictional application.
Fine Tuning	Fine tuning is a machine learning technique used to adapt a pre-trained model to a specific task or domain. It involves taking a model that has already been trained on a large dataset and further training it on a smaller, more specialized dataset relevant to the desired task. This process allows the model to leverage the knowledge gained from the initial training while adapting its parameters to better suit the new task. Fine tuning is commonly used with large language models and other deep learning architectures to improve their performance on specific applications such as sentiment analysis, named entity recognition, or domain-specific language generation.

Generative Artificial Intelligence (GenAI)	Generative AI is that which uses past data to generate novel outputs. A good example of this is ChatGPT, which uses an LLM to accurately predict the next token based on user input. When an LLM can predict the next token accurately enough it is capable of responding in a way that is often indistinguishable from a newly written sentence. This same concept applies to image generation, voice cloning, and other generative AI. It's still doing predictions, but in such a manner that the thing it predicts can appear novel.
Large Language Model (LLM)	A type of artificial intelligence (AI) program that uses deep learning algorithms to recognize, generate, translate, and summarize written language. Notably, LLMs have the ability to achieve general-purpose language understanding and generation.
Machine Learning	A mechanism and technology by which a computer can be trained to use existing data and learn how to perform a specific task, making predictions and improving over time.
Multimodal	The ability of a model to process and understand multiple modes or types of input data, such as text, images, audio, and video.
Natural Language Processing (NLP)	A field of AI that focuses on enabling computers to read, hear, understand, and interpret human language.
Neural Network	A set of algorithms modeled after the structure of the human brain, designed to recognize patterns.

Parameters	The number of individual variables present within a model that are learned during the training phase. They are what store the patterns, insights and relationships obtained from the training data.
Prompt Engineering	Prompt Engineering involves designing and refining the inputs (prompts) given to a LLM to elicit the most effective, accurate, or desired responses. It is a skill that involves understanding how the model interprets and responds to different types of language and structure in prompts. Effective prompt engineering can greatly influence the quality and relevance of the model's outputs, making it a critical aspect of interacting with LLMs.
Retrieval Augmented Generation (RAG)	Retrieval-Augmented Generation is a technique in machine learning that enhances the response generation capabilities of a language model by integrating external information. It combines a pre-trained language model with a retrieval system, which fetches relevant documents or data from a large corpus. This retrieved information is then used by the language model to provide more accurate, detailed, and contextually enriched responses. RAG is particularly useful for tasks requiring up-to-date or detailed knowledge beyond the model's training data. Good examples of RAG systems are Bing Chat from Microsoft and Google’s Search Generative Experience.
Reinforcement Learning	A type of machine learning where an agent learns to make decisions by performing certain actions and receiving rewards or penalties in return. This is commonly used in areas like robotics, gaming, and navigation. It is the technique that powered the famous Go playing model AlphaGo Zero.

Reinforcement Learning with Human Feedback (RLHF)	A technique in machine learning, particularly in the realm of reinforcement learning (RL), where the learning process is guided by feedback from human trainers. The concept combines traditional reinforcement learning, where an agent learns from environmental rewards or punishments, with direct input from humans.
Token	A token refers to the smallest unit of data that a large language model processes. It can be a word, part of a word, or even punctuation. Tokens are used by the model to analyze and generate text, with each token representing a piece of the overall input or output.
Vector Database	A type of database optimized for storing and searching high-dimensional vectors, which are numerical representations of data points. Vector databases are used to store embeddings, which are dense vector representations of entities such as words, sentences, images, or other data types. This is the technology that enables Retrieval Augmented Generation (RAG).

Introduction: Why AI in Advertising?

Artificial intelligence is evolving quickly, with products like ChatGPT being estimated to have reached 100 million active users in just two months[1]. For comparison, the next most recent growth curve was TikTok, which took nine months to reach the same volume. This rapid growth and use in daily and business life is resulting in terminology and concepts being partially defined or incompletely understood. This document aims to set out a foundation of understanding about artificial intelligence to support discussions and documentation about how this technology will affect advertising in the future.

In the business of advertising, two activities are fundamental:

- the scientific pursuit of optimization, and
- the art of creating persuasive media

Artificial Intelligence intersects remarkably with both domains. Intelligence applied to business automates processes, heralding a future where the labor-intensive aspects of digital advertising are significantly reduced. AI not only holds the potential to revolutionize the way we derive insights for marketers and measure performance in a privacy-safe way, but it also paves the way for agencies to explore new creative avenues, and is poised to empower publishers with the ability to craft deeply personalized media experiences.

This document aims to distinguish between the LLMs/foundation models and ad products that are powered by those models. The first part of this document will focus on the use cases for the models themselves. The second part will explore how they're currently being used for advertising use cases.

[1] https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

What is Artificial Intelligence?

Types of Artificial Intelligence

Most conversation today is about Generative AI–this is where rapid enhancement and innovation is commencing. We outline the other types of Artificial Intelligence for context, and to help understand where machine learning, a concept inherent in much of digital advertising fits in.

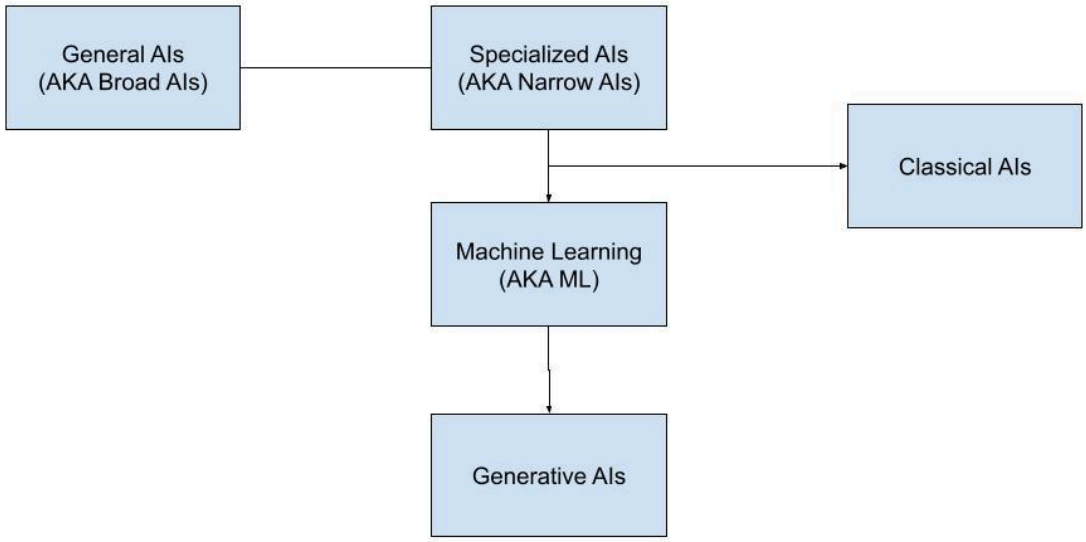

General AIs

Artificial General Intelligence (AGI), alternatively called a Strong Artificial Intelligence, refers to a type of AI that has the ability to understand, learn, adapt, and implement knowledge across a wide range of tasks at a level equal to or beyond human capability. This type of AI could be comparable to what's represented in entertainment such as HAL 9000 in Arthur C. Clark's *Space Odyssey*[2] series or the J.A.R.V.I.S.[3] system in the Marvel Cinematic Universe. Unlike Narrow AI, which is designed to perform specific tasks such as voice recognition or image processing, AGI can theoretically handle any intellectual task that a human being can do.

[2] https://en.wikipedia.org/wiki/2001:_A_Space_Odyssey
[3] https://en.wikipedia.org/wiki/J.A.R.V.I.S.

While the concept of AGI is fascinating and holds immense potential, it's important to note that it currently remains largely theoretical. Despite significant advancements in the field of AI, we still need to develop machines that possess general intelligence comparable to human intelligence. The rest of this document will focus on Narrow AIs as defined below.

Narrow AIs

Narrow Artificial Intelligence, also known as Weak AI or Specialized AI, refers to artificial intelligence systems that are designed and trained for a specific task. Unlike General Artificial Intelligence, Narrow AI does not possess the ability to understand, learn, or apply knowledge beyond its specific function. They do however, have the ability to predict an outcome, with increasing accuracy, based on natural language processing and training sets, thus giving the appearance of learning and improvement.

Examples of Narrow AI are all around us and form part of our everyday lives. From Google's search algorithms to autonomous vehicles, these AI systems excel in their specific tasks but lack the broader understanding and adaptability of General AI. Despite their limitations, Narrow AI systems have brought about significant advancements and efficiencies in various fields including healthcare, finance, transportation, entertainment, and more.

Machine learning

Machine learning (ML) is a subset of Narrow AI that focuses on the development of algorithms and statistical models that enable computer systems to learn from data, identify patterns, and make predictions or decisions without being explicitly programmed.

Machine learning algorithms are designed to learn from data through experience and feedback. They use statistical techniques to identify patterns and relationships in large datasets and then use this knowledge to make predictions or decisions.

Machine learning algorithms can adapt and improve over time as they are exposed to more data. This makes them more flexible and able to handle complex tasks that would be difficult or impossible to program using traditional rule-based systems.

Generative AI

Generative AI is a subset AI that focuses on the development of algorithms and models that enable computer systems to generate new content, such as text, images, music, or videos, based on patterns learned from existing data or content.

Generative AI algorithms are designed to learn from large datasets of existing content and then use this knowledge to generate new, original content that is similar in style, structure, or characteristics to the training data. They use deep learning techniques, chiefly a type of neural network called a transformers network, to identify patterns and relationships in the data and then generate new content based on these patterns.

These algorithms can adapt and improve over time as they are exposed to more data, allowing them to generate increasingly realistic and diverse content. This makes them suitable for tasks that require creativity, such as generating artwork, writing stories, composing music, or designing new products.

We'll give an overview on:

- Large Language Models (LLMs)
- Image Generation & Editing
- Voice Cloning and Audio Generation
- Text to Video Models
- Model Training Techniques

LLM Use Cases

Large language models are where much of the exciting work in generative AI is taking place. When ChatGPT burst onto the scene in late 2022 many quickly learned that LLMs can be used to help polish up an email, fix a line of code, or answer an obscure question. Many also quickly learned that AI would produce convincing falsehoods coined hallucinations. With the release of GPT-4 in March of 2023 it became increasingly clear that this technology was going to be able to assist workers in a variety of daily tasks, not specific to advertising. These models are available commercially and open source. A comprehensive list of open source resources today can be found at huggingface.co[4] Before going into specific advertising applications, we wanted to ensure a base-level understanding of the model capabilities. Here are just some of the opportunities where you can utilize LLMs in your daily work.

1. **Edit, critique, create new written content**: With access to an LLM you have a dedicated editor and co-writer for even the simplest tasks. Ask an LLM to help you check for grammatical errors, rewrite for clarity or tone, or produce an outline or first draft for you to evaluate and build from.

2. **Produce better, more functional code**: Now, every software developer has access to a pair programmer with tools like Github Copilot and ChatGPT. LLMs

[4] https://huggingface.co/

can't outperform the very best coders, but, developers who used coding assistants like Github Copilot were 55% more productive than their peers and reported that using the tool made them feel 75% more fulfilled[5].

3. **Knowledge retrieval**: AI search can allow employees, clients, and customers to ask questions and get up to date information based on your own documentation. This can allow organizations to quickly onboard new employees and deliver answers to customers without taking time away from more pressing tasks.

4. **Streamlining and simplifying the Customer Service experience**: AI agents powered by LLMs can help Customer Service teams answer the same set of repetitive questions, automatically.

5. **Language translation and communication**: Market research and advertisements can be conducted and translated in dozens of languages with these tools.

6. **Note taking**: Note taking AI, like Zoom's IQ, Read.AI, and Fireflies free up employees to focus on making contributions to the meeting.

7. **Structure unstructured data with LLMs**: LLMs rely on a few machine learning concepts to make them work. One of the underlying techniques is textual classification. As a result LLMs are quite adept at taking a large corpus of text and classifying it based on a set of predefined parameters. For example, imagine you have a set of RFPs from advertisers. By defining the categories the RFPs fall under, ie CPG, beauty, tech etc, you can instruct a LLM to iteratively categorize these RFPs.

8. **And many more**: We are only scratching the surface of what can be done with LLMs. As the technology becomes more reliable new opportunities will open up.

Image Generation and Editing Use Cases

Image generators have been publicly available for quite some time, leveraging diffusion models, they are approaching photo-realistic quality in the last two years. In some ways, image generation as a field is more mature than text generation. AI images are competitive with traditional image generation, especially with an experienced graphic designer behind the wheel. Studies suggest that people are only able to detect AI

[5] http://tinyurl.com/iabcopilotstudy

generated images about 53% of the time[6]. AI images, likely sooner than most realize, will be indistinguishable from an image created by a person or taken by a camera.

1. **Image generation**: Image generators have the ability to create photorealistic imagery of almost anything the user can imagine. Stock photography of a woman shopping? A background for your product picture? AI image generators can produce these in moments.

The image on the left is the traditional stock photo taken by Miguel Andres Parra and published for free use by Pexels. The image on the right was created by feeding an AI generated description of the image on the left into MidJourney V6.

2. **Image editing**: Using generative fill a relative novice can easily edit an existing image in a variety of ways. This isn't without limitations, it is not necessarily photorealistic (yet) and struggles with text (for now). But the ability of a novice to easily remove a thumb that is obscuring a camera or add a bicycle helmet to a cat upskills all employees who work with visual imagery instantly.

[6] https://www.nexcess.net/resources/ai-vs-human-study/

A thumb being removed from a photo in just one click in Photoshop.

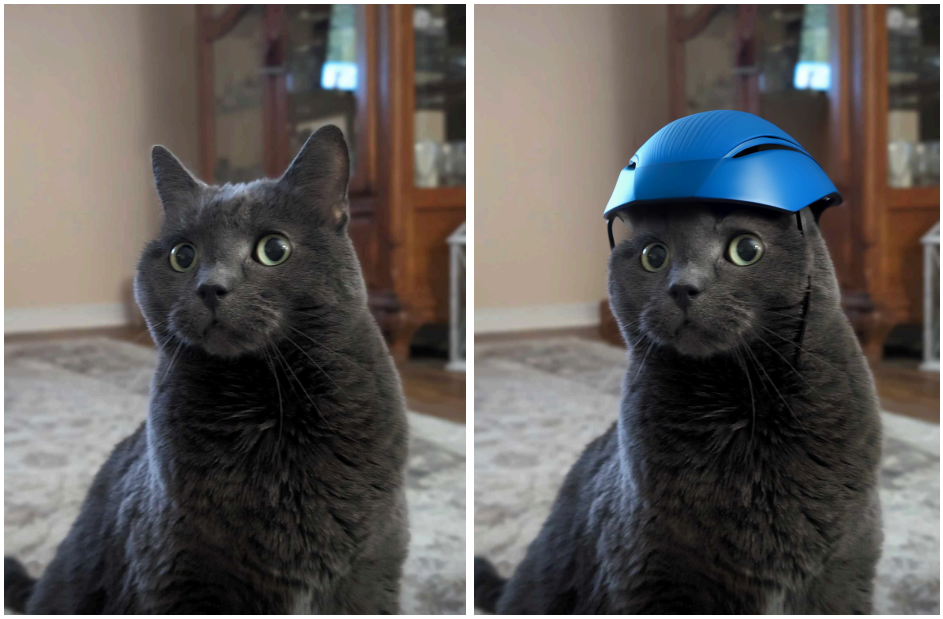

An image of a cat before and after a tiny bicycle helmet has been added to their head using Photoshop's generative fill.

3. **Video editing using image generation**: Technically savvy individuals with sufficient compute can heavily modify videos frame by frame using image editing. The technical barriers will remain high, but AI powered video editing can open up the opportunity to create unique video content. New tools look to push this field even further with end-to-end text-to-video generation likely coming to consumers this year.

Voice Cloning and Audio Generation

Audio generation broke onto the scene with deep faked songs by Drake and The Weekend in 2023[7]. Fully AI generated music has come to fruition in full force this year

[7] https://techcrunch.com/2023/04/17/uh-oh-an-ai-generated-song-by-drake-and-the-weeknd-went-viral/

with upstarts Suno and Udio offering tools that allow users to provide a simple description of the kind of music they want to listen to and receive a full song with vocals in just moments.

On the positive end individuals who have lost their ability to speak can once again speak in their authentic voice just by synthesizing a brief recording, podcasts can hand off ad reads to a robot, and advertisers can quickly translate voice over into dozens of languages to reach new audiences at a lower cost than ever.

Text-to-Video Models

Text-to-Video is the most nascent of the recent generative AI technologies. We've already seen rapid changes in 2024 alone and continue to expect continued evolution in 2025.

Generate video clips: Pika[8], Runway ML[9] Google's Veo[10], Meta's Emu Video [11]and OpenAI's Sora[12] are examples of these types of models. Sora and Veo in particular are stunningly realistic, capable of maintaining consistency, producing realistic reflections, and achieving photorealism on video clips up to a minute long. It won't be long before the ability to generate an entire video ad using only text-to-video models is common, cheap, and effective.

The above examples are far from exhaustive. The real value for advertisers and those in the digital ad supply chain will be finding ways to deploy generative AI both to optimize internal processes and to customize content. Each organization is different and will face different challenges, remaining open, flexible, and thoughtful about your AI use will be key to seeing success with the technology.

Model Optimization Techniques

Pre-trained models

Pre-training models is the process of exposing the model to data so it begins to learn the characteristics present to complete a specific task (such as detecting cats and dogs) or understanding the relationships within the data itself so they can be understood and used to generate similar things later on. Pre-trained models are like off-the-shelf tools that can be used as is or easily adapted to your specific needs.

---

[8] https://pika.art/home
[9] https://runwayml.com/
[10] https://deepmind.google/technologies/veo/
[11] https://emu-video.metademolab.com/
[12] https://openai.com/sora

For example, pre-trained models can be used for tasks such as sentiment analysis, question answering, and text classification. They can also be fine-tuned for specific industries or tasks, such as taking a model that can generate text and specializing it to draft legal briefs making them even more effective.

Fine-tuning

Fine-tuning is a way to take a pre-trained language model and make it better at a specific task or relevant for a particular industry. Think of it as taking a general-purpose tool and customizing it for your specific needs.

For example, let's say you're a marketer who wants to use a language model to generate product descriptions for your e-commerce site. You could fine-tune a pre-trained model to learn the specific language and tone used in your industry or by your brand, so that the generated descriptions are more relevant and engaging.

Or, if you're an ad tech provider, you could fine-tune a pre-trained model to improve its ability to analyze and categorize text data, such as ad copy or social media posts. This would help you provide more accurate and relevant ads to your clients.

RAG

RAG (Retrieval-Augmented Generation) is a technique, enabled by vector databases, that allows you to add context and information to a model's understanding that may not be present in its training data, allowing it to understand new concepts without having to retrain the entire model which is a costly process. It's like having a super-smart assistant that can not only generate text, but also access and use external knowledge to make it even better.

The trade off with RAG though is that the information has to fit into the context window of the model impacting the total prompts or prompt size you may have available. As a marketer, you could use RAG to generate social media posts that are not only engaging based on how the model has been trained, but also informed by the latest external data, such as customer reviews or industry trends. This would help you create more effective and targeted campaigns.

Token Definition

In the context of large language models (LLMs), tokens are fundamental units of text used for processing and generating language, with 1,000 tokens roughly equal to 750 words. Costs are based on the number of tokens processed, with separate charges for input and output tokens. For example, if input tokens are priced at $5.00 per million and output tokens at $15.00 per million, using 2,000 input tokens and generating 1,500 output tokens would cost approximately $0.01 for input and $0.0225 for output, totaling $0.0325. Costs continue to decrease as labs develop more optimization models.

How AI is Applied to Digital Advertising

The complexity of digital advertising has progressed. What once was a simple banner ad served to everyone, is now highly sophisticated, automated systems that connect multiple machine learning powered technologies together in real time. Machine learning has been more and more integrated into all parts of the digital advertising workflow. This increase in prevalence continues to improve the advertiser's ability to present the right creative, to the right person, at the right time to drive the desired business outcome.

Advertisers and Marketers

For advertisers and marketers, AI accelerates creative processes and can inspire new creative outcomes that previously would not have been explored by human intelligence alone. Creative agencies can ideate and iterate much more quickly with generative AI tools. Campaign analysis and optimization could be accelerated by using AI in a similar manner. By leveraging AI tools, these professionals can dissect and interpret vast datasets more efficiently, enabling them to craft strategies that resonate deeply with their target audiences. Moreover, we believe the future of advertising should be built on privacy-safe technology, and AI can help advertisers use fewer signals, while driving the intelligence, measurement and outcomes they need. AI facilitates a more nuanced and individualized understanding of consumer behavior and engagement, leading to more effective, personalized and successful advertising campaigns.

Publishers and Media Owners

In the era of AI-powered media, the notion of "content" becomes incredibly dynamic and malleable, offering publishers and media owners the chance to deliver unique and fully immersive experiences tailored to individual preferences. This shift represents a significant step forward in our evolution of individualized media. In the past, we all consumed the same content, exemplified by the cultural phenomenon of the "Who shot JR?" cliffhanger from the popular 1980s TV series "Dallas," which captivated millions of viewers simultaneously. Today, we have an enormous buffet of choices and recommendations. The future, however, promises a highly personalized AI-media universe where content is not only distributed but also specifically generated to meet each viewer's unique interests and contexts. This "plasticity" of media allows for a custom-made, dynamic experience at the moment of engagement, ensuring content is perfectly suited to every individual's tastes and interests.

Ad Technology companies

The value of ad technology is directly related to the quality of insights it derives from data. AI can produce exponentially more insights from data, both in its ability to appreciate the meaning of data directly and in its ability to be nimble enough to handle vast data points without loss of nuance or fidelity. AI makes insights possible that previously required human experts, and at a speed, over a scale of data, that was never before possible. Specifically, advancements in AI have made it possible to reach audiences most likely to take the desired action, without requiring manual selection of audience segments and behaviors.

The introduction of GenAI in advertising is not merely a new technology; it's a fundamental acceleration of the industry to automation, personalization and preservation of privacy. From delivering customized content to autonomously completing complex data analysis to making effective ad decisions using far less personal data and then writing the code to implement the new learning, AI is set to redefine nearly every aspect of our business in the years to come.

Creative

How is machine learning being used to date?

The personalization of creative, at scale, has historically been a challenge. How to tailor the message of an ad to feel built for a particular person or group required development of many different versions of creatives. Multiple campaign permutations needed to be deployed and managed to granularly segment audiences and guesses needed to be made upfront as to which creatives would perform the best for which audience.

Machine learning first began to help solve this problem through creative optimization. An algorithm could take multiple pieces of creative up front and, as the ads were delivered, invest budget toward the creative performing the best against a given KPI.

These machine learning systems have continued to grow in sophistication They now allow for optimization of creative components[13]. As an example, an advertiser could

[13] https://www.facebook.com/business/news/Introducing-Enhanced-Gen-AI-Features-and-Other-Tools-to-Help-Build-Your-Business

provide 10 different titles, 10 different images, 5 different videos, and 5 different descriptions and a machine learning system can optimize each creative attribute to the audience. This involves finding the people most likely to take the action the advertiser wants and assembling the creative most likely to drive the desired outcome. Adjusting creative elements is only one example of many. Machine learning can also optimize other parts of the ad experience for each user based on their historic preferences. This could include removing or expanding parts of the advert, changing the associated sound or type of music, adding depth of motion, and visual touch ups to enhance the look and feel.

Additionally, with the use of multimodal generative AI technology and the power of LLMs, creative analysis can be done at an unprecedented level in the order of millions of ad units to uncover trends and patterns and to help understand the most impactful aspects of creativity, whether mechanical features, visual/copy features, or thematic features (e.g. distinguishing whether content is humorous and what type of humor resonates best). This will augment and accelerate the power of human creativity, offering powerful insights for creatives to leverage in their ideation process.

Machine learning is helping reduce the investment needed to follow the shift from display to video. Tools are available that can add motion to still images, going beyond previous tricks of panning or zooming. The ability to rapidly edit longer form content to fit shorter mediums is another machine learning use case. Smart tooling can select which parts of the video, and in what order to reuse the existing material.

Even though generative AI is still in very early stages of use in digital advertising, it is available to advertisers. Simple tools to provide basic editing, like removing parts of a creative and infilling, background expansion to change proportion without distortion, and text improvements with alternative copy recommendations are in use today.

Where are the key considerations for the next year?

Existing systems that automate and simplify optimizing creatives will continue to improve and will be further augmented with generative AI features that create and adjust images, text and videos in increasingly sophisticated ways to drive more outcomes.

Generative AI could become more widespread in the design and operational processes as a way to significantly decrease the costs of producing compelling creative. Some example use cases include ideation and rapid prototyping tools to shorten time between concept and first draft. Or systems that are able to translate, with minimal effort, copy,

audio or text to match the styles and framing of an original but translate to different languages. Building on this use case, would also include the ability to reuse existing material but replace or swap out products or other features without having to reshoot everything. We will also likely see experiments with generative AI tools that could have a fine-tuned understanding of the voice and aesthetic of a given brand. These types of tools could be able to generate 'on brand' text, images and eventually even video from scratch.

Content Creation

How is AI being used to date?

The creative section above provides insights into the present and future of AI in ad creative generation. This section will also explore more traditional publisher content creation using AI.

With the explosion of LLMs and image generation models as discussed above, creating content with AI is as simple as a few strokes on a keyboard. Just because it's easy to do, doesn't mean everyone is doing it. As capabilities have evolved, publishers and content creators have been at the forefront of questions about its application to their business. Should we use these tools to create articles for our website? Many organizations are creating their own policies on the matter, reflecting what works for their business and values. An example of this is Thomson Reuters AI Principles[14].

Another use case is in the generation of both music and speech. Like with other mediums it is already possible to generate professional sounding audio clips from a simple text prompt. These can be short pieces of looping background music, bespoke intro and outro music, full songs or even audio samples of everyday objects for use by creators in their content. A similar use case is in speech generation that can take the lines of script, ad copy or podcast conversation and produce professional sounding voice overs and narration. The use of generative AI for audio reduces the cost of creation and increases its accessibility to a much broader audience. It also decreases audio creation and production times by allowing faster prototyping and iterations with creations taking minutes instead of hours or even days.

[14] https://www.thomsonreuters.com/en/artificial-intelligence/ai-principles.html

Where are the key considerations for the next year?

We note more on copyright considerations in the Responsible Use section and further legal analysis has been done for the US Market through IAB's Legal Affairs Council[15]. For brevity, key considerations for the next few years will vary by jurisdiction and may depend on outcomes from court cases contemplating copyright ownership of AI generated content and use of copyrighted content to train AI models.

Targeting/Addressability

How is machine learning being used to date?

Machine learning uses to date for targeting and addressability can largely be grouped into the fields of workflow automation, real time bidding, segmentation and identification.

Workflow automation can be significantly optimized using AI in campaign creation. By integrating targeting and persona information from briefs directly in natural language or semi-structured raw data in spreadsheets, AI can quickly create campaign setups with appropriate segments and demographics. Further automation allows for real-time adjustments and optimizations, ensuring campaigns are constantly tailored to the most promising audience segments based on dynamic data inputs. Furthermore, by allowing flexibility in constraints, AI systems can optimize campaigns faster over larger datasets, enhancing the overall effectiveness and efficiency of ad spend.

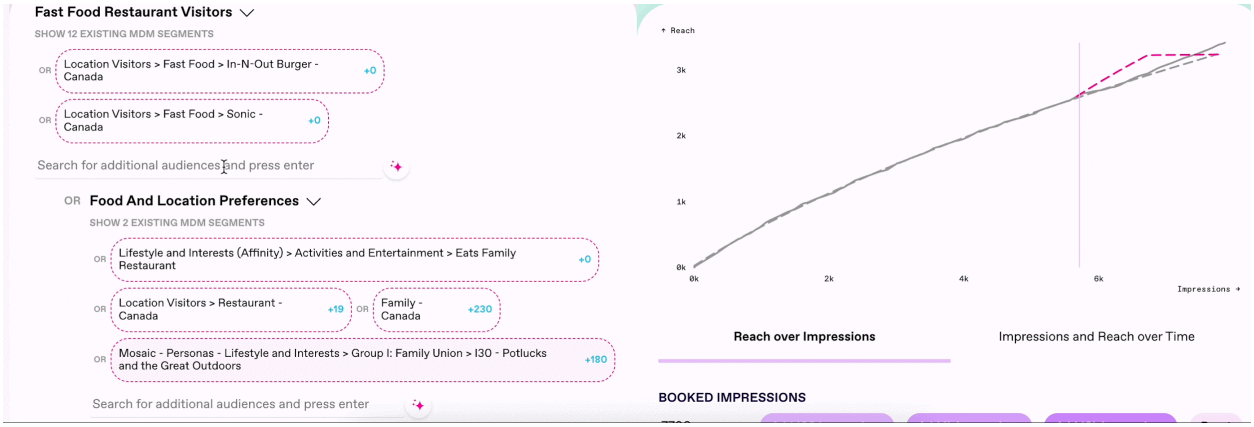

Screen cap of audience discovery and outcome forecasting with AI - source Madhive

15 https://www.iab.com/guidelines/legal-issues-and-business-considerations-when-using-generative-artificial-intelligence-in-digital-advertising/

Real-time bidding (RTB) utilizes machine learning to dynamically adjust bids for ad inventory in low-latency, high variability auction markets, optimizing the likelihood of ad engagement based on the user's current context and historical behavior. By analyzing large amounts of data, including household demographics, browsing habits, and engagement history, machine learning algorithms predict the potential value of an ad impression and adjust bids accordingly. This ensures that advertisers allocate their budgets more effectively, targeting the most promising opportunities in real time, which maximizes return on investment and enhances the relevancy of ads served to users.

Machine-learned audience segmentation allows advertisers and publishers to identify and test targeting which can lead to efficiencies in their business. This technology applies statistical analysis to audience data to uncover patterns that segment the dataset into distinct groups. Each group is characterized by specific behaviors, such as a higher likelihood to purchase an advertiser’s product or greater engagement with a publisher’s content.

Identification ensures consistent user experience across devices, by maintaining and transferring relevant contextual information. For example, if a user browses a car manufacturer’s website on a desktop, they might later see a car ad while playing a mobile game. Identification can be deterministic, such as when a user logs into the same account on multiple devices. Alternatively, it can be probabilistic, relying on data signals — such as device type, location, and browsing patterns — to infer identity. Using the latter method, machine learning algorithms analyze these signals to predict and recognize user identity across different contexts.

Where are the key considerations for the next year?

AI enables new opportunities in targeting and addressability, particularly in the realm of targeting. Advanced language and visual models have made significant strides in understanding nuance, sentiment, and even viewers' intentions. For instance, a car advertiser can leverage a language model to analyze the text of a webpage, determining not only if the page is relevant to car shoppers but also differentiating the intent that brought the user to the page. This allows for targeting based on both contextual relevance and intentional relevance.

As traditional cross site identifiers such as third-party cookies, IDFAs, and IP addresses become more scarce, the importance of AI-powered advertising is growing for publishers and advertisers alike. This technology supports ad placement in environments where individual addressability is limited, ensuring relevance that also aligns with privacy standards and norms. By leveraging the power of AI, advertisers can

increasingly reach their target audiences without relying on cross site user level data, striking a balance between personalization and privacy. Some key areas of development in the next year include:

- New capabilities in precision targeting based on nuance, user intent and context;
- Advances in privacy enhancing ad targeting and measurement techniques;
- Industry standards for best practices for the responsible use of AI in advertising

Measurement and Insights

How is machine learning being used to date?

Machine learning applies to the measurement of impressions by quantifying some of the complex aspects of impression data, for example viewability and anti-fraud. In each of these cases signals from data form the basis, often proprietary, for determining whether an ad view is indeed a billable impression.

Another application of machine learning in measurement has been in response to browser and operating system changes to traditional digital advertising event signals. For example the introduction of Apple's App Tracking Transparency (ATT)[16] framework and the decreasing availability of third party cookies across all web browsers limit the ability to make deterministic connections to attribute events that occurred on an advertisers website or mobile app back to a given advert. To fill this gap advertising systems have invested heavily in probabilistic attribution systems powered by machine learning that are able to estimate, based on the limited deterministic data, how a campaign performed across all users.

Insights encompasses a wide range of valuable reports in advertising, including things like behavioral & cultural trends, competitor analyses, customer journey maps, and any other synthesis of data that allows advertisers and publishers to better achieve their business objectives. In the field of insights, few examples of machine learning techniques are widespread, as to-date this has been a discipline dominated by data science and analysis techniques crafted by human experts with experience and intuition.

[16] https://developer.apple.com/documentation/apptrackingtransparency?language=objc

Where are the key considerations for the next year?

Insights will be particularly accelerated by the adoption of AI. Already the ability to refine, slice and explain data with natural language questions has accelerated the ability for analysts to combine their human ingenuity and contextual understanding of campaign goals over increasingly vast datasets. Going forward we should expect to see more insights derived from the ability of machines to apply step by step reasoning to decompose questions and find answers in data. AI will be able to autonomously test and validate theories about advertising audiences, and report back on those findings. We expect AI to help identify trends from campaign performance datasets, for example derive customer and media consumer intentions or interest from campaign results. These types of insights could then be applied to future targeting and addressability challenges as discussed previously.

Responsible Use Considerations

This document provides a wealth of information about available AI solutions for all parts of the advertising industry. We've highlighted tools for efficiency in day to day business tasks, like writing, developing software and editing images. We've also discussed AI use cases more specific to advertising in the form of ad creative, publisher content, audience targeting and measurement.

There is no shortage of applications of AI to our businesses. Which solutions to enable, where, and how are more of the fundamental questions for advertising professionals. This document won't be able to dictate an answer. Each business will have its own set of questions and consideration depending on its use case. We've outlined a few that are common questions across the myriad of AI applications.

- What are your business and marketing goals and how does AI serve it?
- What are your considerations for organizing data inputs required for an AI model?
- What are your considerations for using the outputs from any AI model?

These questions can be healthy for any organization as they fundamentally require increased coordination between executive teams. Deciding what goes into a model and what comes out of a model creates a healthy dialogue between marketing and technology teams.

Handling Bias - Define Bias for AI

Artificial Intelligence (AI) is only as good as the data it's trained on, and bias can creep in from various sources. In traditional machine learning, bias can occur when:

- Datasets are imbalanced or skewed, leading to unequal predictions or outcomes.
- Features are selected or weighted in a way that unfairly advantages or disadvantages certain groups.
- Models are trained on historical data perpetuating existing biases.
- Reinforcement learning from human feedback (RLHF) where individuals all share similar views and beliefs

To address these issues, consider:

- Data curation: Ensure datasets are diverse, representative, and balanced.
- Data debiasing: Use techniques like pre-processing normalization or re-weighting to mitigate bias.

- Fairness metrics: Monitor and optimize metrics like equality of odds, statistical parity, or counterfactual fairness.
- Diverse fine tuning: Leveraging human feedback from a broad and diverse set of individuals

Generative AI, like text or image generation, introduces new bias concerns:

- Modeled biases: Generative models may reproduce and amplify biases present in the training data.
- Content bias: Generated content may reinforce harmful stereotypes or perpetuate offensive language.
- Lack of diversity: Models may produce homogenous outputs, neglecting underrepresented groups or perspectives.
- Extremities: A model may swing to extremes of a spectrum due to excess of fine tuning and produce content that is overtly biased beyond the data set or real world knowledge

Across the industry, approaches to algorithmic fairness are still evolving, particularly as it relates to digital advertising. But we know we cannot wait for consensus to make progress in addressing important concerns about the potential for discrimination. With that, we need to consider and continue to advocate for:

Diverse and balanced training data, including counter-narratives and diverse perspectives.

- Regular auditing and testing for bias in generated content.
- Implementing mitigation strategies like debiasing techniques, adversarial training, or human oversight.

The overall goal is to ensure the models and their output represent the diversity of the audience. Proactively looking for bias and taking feasible steps to correct it is a continuous process. Algorithmic fairness is an ongoing process that requires evaluation and improvement of AI systems to ensure they serve all individuals and groups equally and without bias.

Copyright

As we started discussing in the Content Creation section above, further questions arise in using AI generated content that takes into consideration how the content was generated. As discussed in the overview sections, all of the models used to produce content require a training data set. Questions then arise about the license grants associated with the assets in the training data. Does leveraging copyrighted content for model training fall under fair use[17]? We won't attempt to answer these questions as

[17] https://www.copyright.gov/title17/92chap1.html#107

many of them are being asked in court cases today, and yet to be decided. They are questions you'll want to stay abreast of in your locale, as regulation and rulings will be different in different jurisdictions.

Responsible AI

AI Regulation

The challenge is not whether certain uses of AI should be regulated, but rather how do we balance the opportunities against the risks in the context of digital advertising. In the public discourse and within governments, there is an often appropriate focus on addressing risks from certain uses of AI. At the same time, AI has been foundational to digital advertising for well over a decade and our experience as an industry shows just how beneficial AI can be for economic growth and opportunity. According to an IAB study[18], the ads-powered internet-economy has been growing 7 times faster than the US average, and has created >7 million jobs over the last 4 years. With AI regulation, we have an opportunity to strengthen our economy, to create jobs, and to drive breakthroughs in health and science. We must focus not only on the harms we want to avoid, but also on the potential we want to achieve.

To strike the right balance between enabling innovation and mitigating possible harms, we have advocated for a risk-based framework, which builds upon sector-specific guidance. What does that mean? High-risk uses or applications of AI technology are becoming subject to more stringent regulatory requirements and government oversight, such as EU AI Act[19]. The use of AI for making a medical diagnosis or for assisting with law enforcement, for example, have a different risk profile than use of AI for digital advertising. For example, under the EU AI Act, digital advertising is not listed under high risk. Different sectors will have important nuances. Industry should help establish standards to recognize those nuances. For example, editing techniques (image resizing, cropping, color or brightening corrections, defect correction, or background edits) are certainly not new for advertising.

IAB and IAB Tech Lab have a proven track record and are well positioned to facilitate the development of standards to support a healthy and sustainable ads ecosystem. Two examples we might consider industry standards are privacy enhancing technologies (PETs), and content labeling of generative AI content.

- Privacy-enhancing technologies can help address concerns about the use of personal data in the context of AI. We need industry, government, and academia

[18] https://www.iab.com/insights/internet-advertising-revenue-report-2024/
[19] https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138-FNL-COR01_EN.pdf

to come together to advance the state of the art and establish standards in this area.

- We are strong supporters of disclosure in appropriate contexts. Several of the large online platforms have announced ads policies requiring election advertisers to disclose the use of synthetic media in election ads. But what about the use of AI in more innocuous contexts? What should be the threshold for disclosure? This is a nuanced area ripe for industry discussion.
- Standard guidance on how to handle personal data when integrating AI models in technology stacks. E.g. memorization/model unlearning should be explored. Possible solutions include encrypting personal data or applying data-level mitigations to address personal data contained therein. Use of public cloud AI models or services should be carefully considered.

AI and Transparency

We advocate for a holistic approach to AI safety, security, and privacy, which recognizes that AI models are deployed as parts of broader systems, and different types of testing and transparency are needed across the AI end-to-end value chain. Given our focus on ads, our discussion of transparency is focused on responsibility within ads product applications, specifically generative AI.

As the line between human-produced and artificially generated content becomes increasingly indistinct, people seek to understand where exactly the division exists. . Digital modification of images is not new and is not necessarily nefarious, but with the advent of Generative AI, the volume of digitally created and modified content is increasing rapidly. It raises the question of when and how to identify something as digitally modified.

Watermarking is often mentioned when discussing the labeling of AI content, but there are different ways to disclose the origin of a piece of synthetic content and different tools such to detect disclosure.

- **Direct Disclosure:** Viewer or listener-facing disclosure directly on the content or user interface.
    - Example: Visible Watermarks, visual labels which may be high prominence (e.g., on-content labeling) or low prominence (e.g., one-click away)
- **Indirect Disclosure:** Attaching into either the content file or content itself (e.g. within the image pixels or audio sound waves) but not obviously visible to user
    - Example: Invisible Watermarks or Metadata

Using both invisible watermarking and metadata improves both the robustness of these invisible markers and helps other platforms identify them.

The invisible watermark is applied with a deep learning model. While it's imperceptible to the human eye, the invisible watermark can be detected with a corresponding model. It's fairly resilient to common image manipulations like cropping, color change (brightness, contrast, etc.), screen shots and more.

There are also other methods such as fingerprinting, which generates unique identifiers (hashes) for content to allow it to be reviewed in a database. Fingerprinting differs from invisible watermark because nothing is added to the content (pixels or file), however, when someone changes the picture, it changes the hash, therefore tracking the changes to content along the way.

Today, no method of detecting GenAI content is 100% comprehensive or accurate, and we rely upon advertiser self-disclosure. Industry associations have a helpful role to play advancing open and standard methods for tracing.

When is it appropriate to disclose to the viewers of an ad that the content in the ad is AI generated or enhanced? Concerned about deception from AI content, many users, policymakers, press, and advocacy groups are emphasizing the need for user-facing labels to identify generative AI content. But as noted earlier, not all AI content is deceptive, and not all deceptive content is made with generative AI. An overemphasis on whether an ad is generated may imply that all generated content is deceptive, while all other content is trustworthy. User-facing labels on certain types of content may be helpful in some situations, (e.g., on election content where there is an expectation of unaltered content) they are one of many tools to reduce deceptive content. Many platforms leverage AI powered technologies to reduce and remove harmful or misleading content and to identify which content should be flagged to human reviewers often before the content has been reported by a consumer.

It is important to note that the distinctions between content that is AI generated versus those that are not is ambiguous and getting more so. For example, photo retouching tools rely on AI systems to denoise and improve images as they come out of a smartphone's camera. Labeling any image not purely reliant on the optics of a camera would result in more public confusion and there has yet to be industry alignment on any threshold approach.

Moreover, it is important to avoid unintended consequences that can make labeling requirements counterproductive. Research shows that inconsistent labeling has the

potential to increase trust in unlabeled content. This phenomenon, which is well documented in academic research (see e.g. here[20], here[21]), creates a danger of distorting users' understanding in general, where they will believe all unlabeled content is "real" and may lose trust in the sources of labeled content. Hence, in the near term an overbroad standard requiring labels could punish legitimate advertisers, and potentially aid bad actors use of generative AI. Over time, as we see greater adoption of generative AI, we could face a different problem. There are risks of “notification fatigue” or “banner blindness,” where users begin to ignore labels as they become more pervasive. We’ve seen this with website operators’ attempts to comply with the EU’s General Data Protection Regulation’s website cookie disclosure requirements, which showed that overly broad disclosure obligations can be highly disruptive while providing limited safety benefits.

We believe that industry has a role to play in developing guidelines around AI and ads transparency. A responsible approach to labeling needs to mitigate the risks of labeling and allow for evolution over time.

AI and Content Moderation

We recognize that AI has introduced challenges to ads safety. Good advertisers are using generative AI to optimize their campaigns and increasingly to generate compelling ad creative. However, the same technology can be used by bad actors to create better looking ads with greater reach on faster timelines.

However, AI is also one of the most effective tools to moderate content at scale. For example, Meta open-sourced an AI-powered content moderation software called the Hasher-Matcher Actioner[22] which is designed to apply a digital fingerprint to a piece of harmful content in order to make it easier to spot replicas and pull them down en masse—not just on one platform—but on others as well. The more companies that use the tool and participate in the hash sharing database, the more effective and comprehensive it is at stopping the spread of violating content on the internet.

Online platforms have long used machine learning to enforce their policies at scale. But, while still highly sophisticated, these machine learning models have historically needed to be trained extensively — they often rely on hundreds of thousands, if not millions of examples of violative content. LLMs, on the other hand, are able to rapidly review and

[20] https://psycnet.apa.org/record/2021-17032-001
[21] https://www.researchgate.net/publication/321887941_The_Implied_Truth_Effect_Attaching_Warnings_to_a_Subset_of_Fake_News_Headlines_Increases_Perceived_Accuracy_of_Headlines_Without_Warnings
[22] https://about.fb.com/news/2022/12/meta-launches-new-content-moderation-tool/

interpret content at a high volume, while also capturing important nuances within that content. These advanced reasoning capabilities have already resulted in larger-scale and more precise enforcement decisions on some of our more complex policies.

Ads and publisher platforms should look to leverage the power of LLMs in their ads safety and enforcement efforts, and regulation should enable this important use.

Conclusion

Ever-widening user adoption and continuing breakthroughs make the present an exciting time for Artificial Intelligence. The high tempo of advancements in generative artificial intelligence has garnered a large amount of attention in the press and in the digital media and advertising ecosystems as various organizations and individuals attempt to leverage it to address a number of existing and emerging challenges. As of mid 2024, many of these attempts are still in the experimental phase as practitioners and AI innovators collaborate to discover how these new technologies can be used to augment or disrupt the systems and processes that dominate the current digital advertising landscape.

Both participants and observers of these experiments should endeavor to maintain a mind both open to new ways of working and skeptical of wild claims without supporting evidence that can be independently reproduced and confirmed.
We expect the coming months/year to:

1. Gain a better understanding of how new generative AI tools will change elements of the current digital advertising ecosystem by either improving existing processes and practices or disrupting existing systems and entities.
2. Allow experts to better predict how the addition of expanded computational resources will affect the effectiveness of proposed generative AI digital advertising solutions.

As outlined, we believe there is opportunity for standardization to help fuel the growth of these technologies in digital advertising. We look forward to continued collaboration with the industry.